

[http://www.wvpi.com/index.php?option=com\\_content&view=article&id=15595:accelerating-san-storage-with-server-flash-caching&catid=334:feature-articles&Itemid=2701754](http://www.wvpi.com/index.php?option=com_content&view=article&id=15595:accelerating-san-storage-with-server-flash-caching&catid=334:feature-articles&Itemid=2701754)

## Accelerating SAN Storage with Server Flash Caching

**By Tony Afshary**

THURSDAY, 31 JANUARY 2013

The data deluge, with its relentless increase in the volume and velocity of data, has brought renewed focus on an old problem: the enormous performance gap that exists in input and output (I/O) operations between a server's memory and disk storage. I/O takes a mere 100 nanoseconds for information stored in a server's memory, whereas I/O to a hard disk drive (HDD) takes about 10 milliseconds — a difference of five orders of magnitude that is having a profound adverse impact on application performance and response times.

The lower bandwidth and higher latency in a storage area network (SAN) or network-attached storage (NAS) combine to exacerbate the performance problem, which gets even worse with the frequent traffic congestion on the intervening Fibre Channel (FC), FC over Ethernet, iSCSI or Ethernet network. This storage bottleneck has grown over the years as the increase in drive capacities has outstripped the decrease in latency of faster-spinning drives. As a result, the performance limitations of most applications have become tied to latency more than bandwidth or I/Os per second (IOPs), and this trend is expected accelerate as the amount of data being created continues to grow between 30 and 50 percent per year.

It is instructive to look at the situation from another perspective. The past three decades have witnessed a 3000 times increase in network bandwidth, while network latency has been reduced by only about 30 times. During the same period, the gains in processor performance, disk capacity and memory capacity have also been similarly eclipsed by the relatively modest reduction in latency.

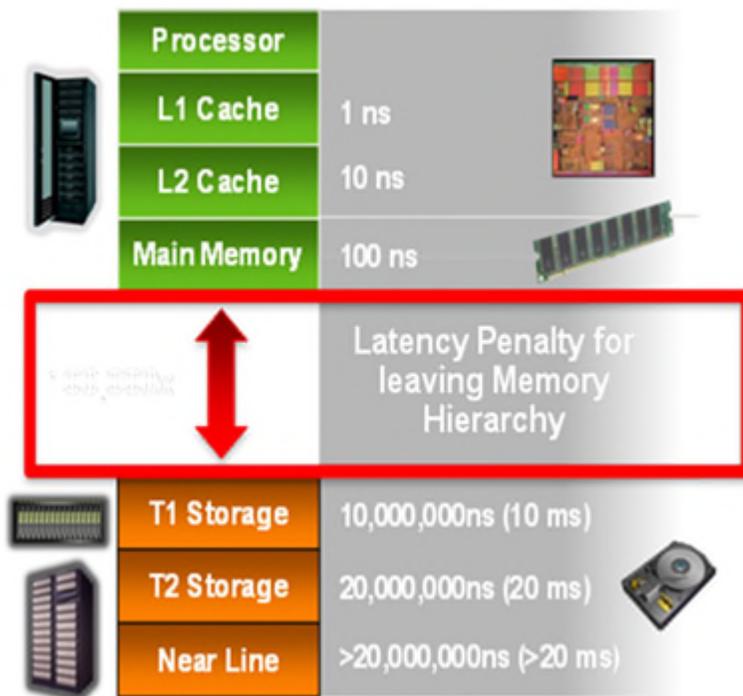
The extent of the problem became apparent in a recent survey conducted by LSI of 412 European datacenter managers. The results revealed that while 93 percent acknowledge the critical importance of optimizing application performance, a full 75 percent do not feel they are achieving the desired results. Not surprisingly, 70 percent of the survey respondents cited storage I/O as the single biggest bottleneck in the datacenter today.

The challenge will only get greater, caused by what LSI calls the data deluge gap — the disparity between the 30 to 50 percent annual growth in storage capacity requirements and the 5 to 7 percent annual increase in IT budgets. The net effect is that data is growing faster than the IT infrastructure investment required to store, transmit, analyze and manage it. The result is that IT departments and datacenter managers are under increasing pressure to find smarter ways to bridge the data deluge gap and improve performance.

## Cache in a Flash

Caching content to memory in a server or in a SAN on a Dynamic RAM (DRAM) cache appliance is a proven technique for improving storage performance by reducing latency, and thereby improving application-level performance. But because the amount of memory possible in a server or cache appliance (measured in gigabytes) is only a small fraction of the capacity of even a single hard disk drive (measured in terabytes) performance gains from this traditional form of caching are becoming increasingly insufficient to overcome the challenges of the data deluge gap.

NAND flash memory technology breaks through the cache size limitation imposed by traditional memory to again make caching the most effective and cost-effective means for accelerating application performance. As shown in the diagram, NAND flash memory fills the significant void between main memory and Tier 1 storage in both capacity and latency.



***Flash memory fills the void in both latency and capacity between main memory and fast-spinning hard disk drives.***

Solid state memory typically delivers the highest performance gains when the flash cache acceleration card is placed directly in the server on the PCI Express (PCIe) bus. Embedded or host-based intelligent caching software is used to place “hot data” (the most frequently accessed data) in the low-latency flash storage, where data is accessed up to 200 times faster than with a Tier 1 HDD, where less frequently accessed data is stored.

Astute readers may be questioning how flash cache, with a latency 100 times higher than DRAM, can outperform traditional caching systems. There are two reasons for this. The first is the significantly higher capacity of flash memory, which dramatically increases the “hit rate” of the cache. Indeed, with some of these flash cache cards now supporting multiple terabytes of solid state storage, there is often sufficient capacity to store entire databases or other datasets as “hot data.”

The second reason involves the location of the flash cache: directly in the server on the high-speed PCIe bus. With no internal or external connections, or no intervening network that is also subject to frequent congestion, the “hot data” is accessible in a flash (pun intended) and in a deterministic manner under all circumstances.

Tests show that the performance gains of server-side flash-based caching are both consistent and significant under real-world conditions. Tests performed by LSI using Quest Benchmark Factory software and audited by the Transaction Performance Council, clearly demonstrate how a PCIe-based flash acceleration card can improve database application-level performance by a conservative 5 to 10 times compared to either direct-attached storage (DAS) or a SAN.

### ***More and Better Flash***

As the pricing of flash memory continues to drop and its performance continues to improve, flash memory will become more prevalent throughout the datacenter. Will flash-based solid state drives (SSDs) ever replace hard disk drives? No, at least in the foreseeable future. HDDs have enormous advantages in storage capacity and in the cost of that capacity on a per-gigabyte basis. And because the vast majority of data in most organizations is only rarely accessed, the higher latency of HDDs is normally of little consequence — especially if this “dusty data” can become “hot data” in a PCIe flash cache accelerator on those rare occasions when it is needed.

The key to making continued improvements in flash price/performance — comparable to that of processors according to Moore’s Law — is advancements in the flash controllers that facilitate ever-shrinking NAND memory geometries, already under 20 nanometers. The latest generation of flash controllers offers sophisticated wear-leveling to improve flash memory endurance, and enhanced error correction algorithms to improve reliability with RAID-like data protection.

These advances are making it possible for PCIe-based flash caching solutions to provide advanced capabilities beyond those available with traditional caching. For example, caching has historically been a read-only technology, but RAID-like data protection for writes to flash memory has the effect of making the cache the equivalent of a fast storage tier. The addition of acceleration for writes to flash cache (which are then persisted to RAID-based DAS or SAN) can improve application-level performance by up to 30 times compared to HDD-only storage systems.

### ***The Future of Flash***

Flash memory has already become the primary storage in tablets and ultrabooks, and a growing number of laptop computers. Solid state drives are replacing or supplementing hard disk drives in desktop computers and the direct-attached storage in servers, while SSD storage tiers are growing larger in SAN and NAS configurations. And the use of PCIe-based acceleration adapters is growing rapidly owing to their ability to bridge the data deluge gap better than any other alternative.

Some of the other advantages of flash (not discussed here) are giving these trends additional momentum. Flash has a higher density than hard disk drives, enabling more storage in a smaller space. Flash also consumes less power, and therefore, requires less cooling. These advantages are equally beneficial at both a small scale in a tablet and a large scale in a datacenter.

Even as flash memory becomes more pervasive throughout datacenters, there will continue to be a need for PCIe flash acceleration cards in servers for quite some time. Indeed, the flash cache is expected to remain the most effective and cost-effective way to accelerate application performance for the foreseeable future.

*Tony Afshary is the director of marketing for the Accelerated Solutions Division of [LSI Corporation](#).*